

# Optional Advanced Tuning for Linux

## Category: Network

This document describes additional TCP settings that can be tuned on high-performance Linux systems. This is intended for 10-Gigabit hosts, but can also be applied to 1-Gigabit hosts. The following steps should be taken in addition to the steps outlined in [TCP Performance Tuning for WAN transfers](#).

Configure the following `/etc/sysctl.conf` settings for faster TCP

1. Set maximum TCP window sizes to 12 megabytes:

```
net.core.rmem_max = 11960320
net.core.wmem_max = 11960320
```

2. Set minimum, default, and maximum TCP buffer limits:

```
net.ipv4.tcp_rmem = 4096 524288 11960320
net.ipv4.tcp_wmem = 4096 524288 11960320
```

3. Set maximum network input buffer queue length:

```
net.core.netdev_max_backlog = 30000
```

4. Disable caching of TCP congestion state (Linux Kernel version 2.6 *only*). Fixes a bug in some Linux stacks:

```
net.ipv4.tcp_no_metrics_save = 1
```

5. Use the BIC TCP congestion control algorithm instead of the TCP Reno algorithm (Linux Kernel versions 2.6.8 to 2.6.18):

```
net.ipv4.tcp_congestion_control = bic
```

6. Use the CUBIC TCP congestion control algorithm instead of the TCP Reno algorithm (Linux Kernel versions 2.6.18 and newer):

```
net.ipv4.tcp_congestion_control = cubic
```

7. Set the following to 1 (should default to 1 on most systems):

```
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
```

A reboot will be needed for changes to `/etc/sysctl.conf` to take effect, or you can attempt to reload `sysctl` settings (as root) with `sysctl -p`.

For additional information visit the [Energy Science Network website](#).

If you have a 10-Gb system or if you follow these steps and are still getting less than your expected throughput, please contact NAS Control Room staff at [support@nas.nasa.gov](mailto:support@nas.nasa.gov), and we will work with you on tuning your system to optimize file transfers.

---

Article ID: 138

Last updated: 02 Oct, 2012

Data Storage & Transfer -> Storage Components -> Network -> Optional Advanced Tuning for Linux

<http://www.nas.nasa.gov/hecc/support/kb/entry/138/?ajax=1>